

SHARP: a distributed, GPU-based ptychographic solver

Stefano Marchesini,¹ Hari Krishnan,¹ Benedikt Daurer,² Filipe R.N.C. Maia,² David A. Shapiro,¹ Talita Perciano,¹ and James A. Sethian¹

¹*Lawrence Berkeley National Laboratory, Berkeley, CA, USA*

²*Uppsala University, Uppsala, Sweden*

(Dated: January 7, 2016)

With ever brighter light sources, fast parallel detectors, and advances in phase retrieval methods, one can process experimental “imaging by diffraction” data to produce sharp images. Additionally, augmenting these technologies with high-throughput processing can enable close to real-time feedback to microscopists. Together, these advances allow rapid images across a range of macroscopic dimensions, through different energies, and time evolution at x-ray wavelength resolution. Here, we introduce set of algorithmic and computational methodologies, packaged as an open source software suite, aimed at throughput and resolution for the coming diffraction limited light source era.

I. INTRODUCTION

The reconstruction of a scattering potential from measurements of scattered intensity in the far-field has occupied scientists and applied mathematicians for over a century, and arises in fields as varied as optics^{1,2}, astronomy³, X-ray crystallography⁴, tomographic imaging⁵, holography^{6,7}, electron microscopy⁸ and particle scattering generally. Numerous experimental techniques that employ forms of interferometric/holographic^{6,7} measurements, gratings⁹, and other phase mechanisms like random phase masks, sparsity structure, etc^{10–17} to help overcome the problem of phase-less measurements have been proposed over the years^{18–20}.

Progress has been made in solving the phase problem for a single diffraction pattern recorded from a non-periodic object^{21–25}. Such methods, referred to as coherent diffractive imaging (CDI), attempt to recover the complete complex-valued wave scattered from the object, giving phase contrast and a way to overcome depth-of-focus limitations of regular optical systems.

More recently an experimental technique has emerged that enables to image what no-one was able to see before: macroscopic specimens in 3D at wavelength (i.e. potentially atomic) resolution, with chemical state specificity. Initially proposed in 1969^{27–31} to improve the resolution in electron or x-ray microscopy by combining microscopy with scattering measurements, ptychography has shown to be a remarkably robust technique for the characterization of nano materials^{27,32–54}.

Ptychography enables one to build up very large images at wavelength resolution (i.e. potentially atomic) by combining the large field of view of a high precision scanning microscope system with the resolution enabled by diffraction measurements. Experimentally, ptychography works by retrofitting a scanning microscope with a parallel detector. In a scanning microscope, a small beam is focused onto the sample via a lens, and the transmission is measured in a single- element detector. The image is built up by plotting the transmission as a function of the sample position as it is rastered across

the beam, and the resolution of the image is given by the beam size. In ptychography, one replaces the single element detector with a two-dimensional array detector such as a CCD and measures the intensity distribution at many scattering angles, much like a radar detector system for the microscopic world. Each recorded diffraction pattern contains short-spatial Fourier frequency information about features that are smaller than the x-ray beam-size, enabling higher resolution. At short wavelengths, however, it is only possible to measure the intensity of the diffracted light. To reconstruct an image of the object, one needs to retrieve the phase, made even more challenging in the presence of noise, experimental uncertainties, and perturbations of the experimental geometry. Since the reconstruction of ptychographic data is a non-linear problem, there are still many open questions²⁶, nevertheless the phase retrieval problem is made tractable by recording multiple diffraction patterns from the same region of the object, compensating phase-less information with a redundant set of measurements. Data redundancy enables one to handle experimental uncertainties as well. Methods to work with unknown illuminations or “lens” were proposed^{35–37,39,55}. They are now used to calibrate high quality x-ray optics^{40–42}, x-ray lasers⁵⁶ and space telescopes⁵⁷. Several strategies, such as Alternating Directions⁵⁸, projections, gradient, conjugate gradient, Newton⁵⁹, spectral methods^{26,49,60}, and Monte-carlo⁴³, have been proposed to handle situations when both sample and illumination function, positions, are unknown parameters in high dimensions, and to handle experimental situations such as partial coherence, background, averaging during flying scans, vibrations^{43–47,59–66}

Here we describe an algorithm approach and open software suite “SHARP” (Scalable Heterogeneous Adaptive Real-time Ptychography) that enables high throughput streaming analysis using computationally efficient phase retrieval algorithms and distributed computational back-end. The high performance computational back-end is hidden from the microscopist, but can be accessed and adapted to particular needs by using a python interface or by modifying the source code.

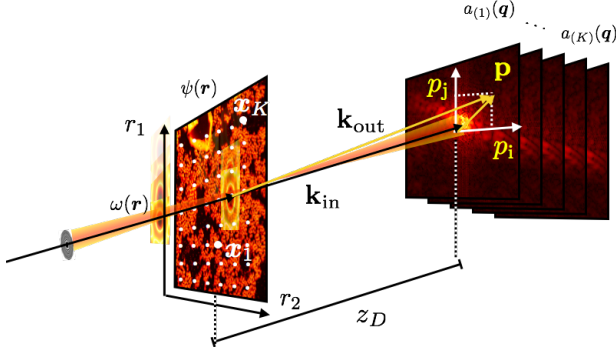


FIG. 1. Experimental geometry in ptychography: an unknown sample with transmission $\psi(\mathbf{r})$ is rastered through an illuminating beam $w(\mathbf{r})$, and a sequence of diffraction measurements $I_{(i)} = |\mathbf{a}_{(i)}|^2$ are recorded on an area detector as the sample is rastered around. Each set of pixel coordinates on a detector placed at a distance z_D from the sample is $\mathbf{p}_{i,j,z_D} = (p_i\hat{x} + p_j\hat{y} + z_D\hat{z})$, where $p_{i,j}$ denote the pixel coordinates and \hat{x}, \hat{y} the axes of the detector. The detector coordinates \mathbf{p}_{i,j,z_D} are related to coordinates $\mathbf{q}_{i,j}$ by $\mathbf{q}_{i,j} = \mathbf{k}_{\text{out}} - \mathbf{k}_{\text{in}} = \frac{1}{\lambda} \left(\frac{(p_i, p_j, z_D)}{\sqrt{p_i^2 + p_j^2 + z_D^2}} - (0, 0, 1) \right)$, $\simeq \frac{1}{\lambda z_D} (p_i, p_j, 0)$, \mathbf{k}_{in} and \mathbf{k}_{out} are the incident and scattered wave vectors that satisfy $|\mathbf{k}_{\text{in}}| = |\mathbf{k}_{\text{out}}| = k = 1/\lambda$, where λ is the wavelength of the x-ray. The direction of \mathbf{k}_{in} and \mathbf{k}_{out} are the same as the direction of the incident beam $\mathbf{k}_{\text{in}} = (0, 0, k)$ and the outgoing beam $\mathbf{k}_{\text{out}} = k\hat{\mathbf{p}}_{i,j,z_D}$. The diffracted photon flux I (photons/pulse/pixel) produced by a sample illuminated by x-rays is given by $I_{\text{data}}(\mathbf{q}) = J_o r_e^2 P \Delta\Omega |\mathcal{F}_{\mathbf{q} \leftarrow \mathbf{r}} E(\mathbf{r})|^2$, where J_o is the incident photon flux density (photons/pulse/area), r_e^2 is the electron cross section, P is a polarization multiplicative factor, $\Delta\Omega$ is the solid angle subtended by a detector pixel at the sample, $E(\mathbf{q})$ is a complex scalar wavefield transmitted by the sample, and \mathcal{F} is a Fourier transform operation. The point-wise product between illuminating function and sample, $\mathbf{z}_{(i)}(\mathbf{r}) := w(\mathbf{r})\psi(\mathbf{r} + \mathbf{x}_{(i)})$, is related to the measurement by a Fourier magnitude relationship $I_{(i)} = \mathbf{a}_{(i)}^2 = |F\mathbf{z}_{(i)}|^2$.

An intuitive graphical user interface provides visual feedback of the diffraction data and reconstructed images, while data is added to the solver or as iterations progress over time. Our software is modular and provides several interfaces, remote visualization, and the file I/O needed to handle the needs of different microscopes, off-line analysis, benchmarking or end-to-end experimental simulations. We use a standard cxi file format⁶⁷ for file input output operations, and standard open source libraries and formats for streaming frames and images across multiple modules.

II. ALGORITHMS

In a ptychography experiment (see Fig. 1), one performs a series of diffraction measurement as a sample is rastered across an x-ray beam. The illumination is formed by an x-ray optic such as a zone-plate. The mea-

surement is performed by briefly exposing an area detector such as a CCD which discretizes the charges accumulated over a square pixel.

In a discrete setting, a two-dimensional small beam with distribution $w(\mathbf{r})$ of dimension $m_x \times m_y$ illuminates a subregion centered at $\mathbf{x}_{(i)}$ (referred to as frame) of an unknown object of interest ψ of dimension $n_x \times n_y$. Here $0 < m < n$, $i = 1, \dots, k$, k is the number of frames, $r = \frac{\lambda p_m}{z_d}$ is a lengthscale (the diffraction limited resolution of the microscope, see Fig. 1 for the definition), and:

$$\begin{aligned} \mathbf{r} &= (r\mu, r\nu), \quad \mu, \nu \in \{0, \dots, m-1\}, \\ \mathbf{x}_{(i)} &= (r\mu', r\nu'), \quad \mu', \nu' \in \{0, \dots, n-m\}. \end{aligned}$$

As $\mathbf{x}_{(i)}$ is rastered on a typically coarser grid, $\mathbf{r} + \mathbf{x}_{(i)}$ spans a finer grid of dimension $n \times n$. Here for simplicity we consider square matrices. Generalization to non-square matrices is straightforward but requires more indices and complicates notation.

In other words, we assume that a sequence of k diffraction intensity patterns $\mathcal{I}_{(i)}(\mathbf{q})$ are collected as the position of the object is rastered on the position $\mathbf{x}_{(i)}$. The simple transform $\mathbf{a}_{(i)} = \sqrt{\mathcal{I}_{(i)}(\mathbf{q})}$ is a variance stabilizing transform for Poisson noise^{68,69}. The relationship among the amplitude $a_{(i)}$, the illumination function w and an unknown object ψ to be estimated can be expressed as follows:

$$\begin{aligned} a_{(i)}(\mathbf{q}) &= \left| \mathcal{F}w(\mathbf{r})\hat{\psi}(\mathbf{r} + \mathbf{x}_{(i)}) \right|, \\ (\mathcal{F}f)(\mathbf{q}) &= \sum_{\mathbf{r}} e^{i\mathbf{q} \cdot \mathbf{r}} f(\mathbf{r}), \\ \mathbf{q} &= \left(\frac{2\pi}{mr} \mu, \frac{2\pi}{mr} \nu \right), \quad \mu, \nu \in \{0, \dots, m-1\}. \end{aligned} \quad (1)$$

where the sum over \mathbf{r} is given on all the indices $m \times m$ of \mathbf{r} , and \mathcal{F} is the two-dimensional discrete Fourier transform. The relationship in Fig. (1) relates the pixel coordinates on the detector (p_i, p_j, z_D) , the sampling coordinate \mathbf{q} , and the length-scale r that maps the sample positions \mathbf{x} to coordinates on a discretized image ψ . The Fourier transform relationship is valid under Far Field and paraxial approximation, and small illumination ("oversampling" condition), which is the focus of the current release of SHARP. For experimental geometries such as Near Field, "Fresnel", Fourier ptychography, through-focus, partially coherent, multiplexed, geometries^{50,52-54} one can substitute the simple Fourier transform with the appropriate propagator⁶⁰. Weighting factors to account for noisy pixel will be added in a future release⁵⁹.

We denote the illumination operator $Q_{(i)}$, $i = 1, 2, \dots, k$, associated with $\mathbf{x}_{(i)}$ that extracts a frame $z_{(i)}$ out of ψ , and scales the frame point-wise by the illumination function $w(\mathbf{r})$ (see Fig. 1):

$$\begin{aligned} Q_{(i)}[\psi](\mathbf{r}) &= w(\mathbf{r})\psi(\mathbf{r} + \mathbf{x}_{(i)}), \\ &= w(\mathbf{r})T_{(i)}[\psi](\mathbf{r}), \\ &= z_{(i)}(\mathbf{r}) \end{aligned}$$

where $z_{(i)}$ represents the frames extracted from ψ and multiplied by the illumination $w(\mathbf{r})$, and $T_{(i)}$ is the operator that extracts a frame out of an image. With this operator, the relationship between an unknown discretized object ψ , and the diffraction measurements $\mathbf{a}_{(i)}$ collected in a ptychography experiment (see figure 1) can be represented compactly as:

$$\mathbf{a} = |\mathbf{F}\mathbf{Q}\psi^\vee|, \text{ or } \begin{cases} \mathbf{a} = |\mathbf{F}\mathbf{z}|, \\ \mathbf{z} = \mathbf{Q}\psi^\vee, \end{cases} \quad (2)$$

Eq. (2) can be expressed as:

$$\begin{bmatrix} \mathbf{a}_{(1)} \\ \vdots \\ \mathbf{a}_{(K)} \end{bmatrix} = \begin{bmatrix} \mathbf{F} & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \mathbf{F} \end{bmatrix} \begin{bmatrix} \mathbf{z}_{(1)} \\ \vdots \\ \mathbf{z}_{(K)} \end{bmatrix}, \quad (3)$$

$$\begin{bmatrix} \mathbf{z}_{(1)} \\ \vdots \\ \mathbf{z}_{(K)} \end{bmatrix} = \begin{bmatrix} \text{diag}(w)T_{(1)} \\ \vdots \\ \text{diag}(w)T_{(K)} \end{bmatrix} \begin{bmatrix} \psi_1 \\ \vdots \\ \psi_{n^2} \end{bmatrix}. \quad (4)$$

where \mathbf{z} are K frames extracted from the object ψ and multiplied by the illumination function w , and \mathbf{F} is the associated 2D DFT matrix when we write everything in the stacked form⁶⁰. When both the sample and the illumination are unknown, we can express the relationship (Eq. 4) between the image ψ , the illumination w , and the frames \mathbf{z} in two forms:

$$\mathbf{z} = \mathbf{Q}\psi = \text{diag}(\mathbf{S}w)\mathbf{T}\psi = \text{diag}(\mathbf{T}\psi)\mathbf{S}w \quad (5)$$

where $\mathbf{S} \in \mathbb{R}^{Km^2 \times m^2}$ denotes the operator that replicates the illumination w into K stack of frames, since $\mathbf{Q}\psi = \text{diag}(\mathbf{S}w)\mathbf{T}\psi$ is the entry-wise product of $\mathbf{T}\psi$ and $\mathbf{S}w$. Eq. (5) can be used to find ψ or w from \mathbf{z} and the other variable.

A. Phase retrieval

A typical reconstruction on SHARP uses the following sequence:

1. Input data $\mathcal{I}(\mathbf{q})$, translations \mathbf{x} . Optional inputs: initial image $\psi^{(0)}$, illumination $w^{(0)}$, and other constraints such as illumination Fourier mask.
2. Initialize illumination, image, and build up \mathbf{Q} , \mathbf{Q}^* , and $(\mathbf{Q}^*\mathbf{Q})^{-1}$, and frames $\mathbf{z}^{(0)} = \mathbf{Q}\psi^{(0)}$;
3. Update the frames \mathbf{z} according to⁷⁰ using projector operators defined in (Eqs. (6,7)) below:

$$\mathbf{z}^{(l)} := [2\beta P_{\mathbf{Q}}P_{\mathbf{a}} + (1 - 2\beta)\beta P_{\mathbf{a}} + \beta(P_{\mathbf{Q}} - I)]\mathbf{z}^{(l-1)},$$

where $\beta \in (0.5, 1]$ is a scalar factor set by the user.

4. Update illumination $w^{(\ell)}$, and image $\psi^{(\ell)}$ using (Eqs. (9, 8)), and additional constraints such as a Fourier low pass or band pass binary filter, or a real space mask.
5. Optional: optimize static background and remove it in the iteration as described in⁶⁰
6. Iterate to 3 until maximum iteration or target is achieved, and return $\psi^{(\ell)}$ and w .

Projection operators form the basis of every iterative projection and projected gradient algorithms and are implemented in SHARP and accessible through a library. The projection $P_{\mathbf{a}}$ ensures that the frames \mathbf{z} match the experiment, that is, they satisfy Eq. (3), and is referred to as data projector:

$$P_{\mathbf{a}}\mathbf{z} = \mathbf{F}^* \frac{\mathbf{F}\mathbf{z}}{|\mathbf{F}\mathbf{z}|} \mathbf{a} \quad (6)$$

while the projection $P_{\mathbf{Q}}$ onto the range of \mathbf{Q} (see Fig. 2):

$$P_{\mathbf{Q}} = \mathbf{Q}(\mathbf{Q}^*\mathbf{Q})^{-1}\mathbf{Q}^* \quad (7)$$

ensures that overlapping frames \mathbf{z} are consistent with each other and satisfy Eq. (4).

The projector $P_{\mathbf{a}}$ is relatively robust to Poisson noise⁶⁸, but weighting factors to account for noisy pixels can be easily added⁵⁹.

Using relationship (5), we can update the image ψ from w and frames \mathbf{z} :

$$\psi \leftarrow \frac{\mathbf{Q}^*\mathbf{z}}{\mathbf{Q}^*\mathbf{Q}} \quad (8)$$

or the illumination w from an image ψ and frames \mathbf{z} ^{36,37}:

$$w \leftarrow \frac{\mathbf{S}^* \text{diag}(\mathbf{T}\psi)\mathbf{z}^{(l)}}{\mathbf{S}^*\mathbf{T}|\psi|^2}, \quad (9)$$

The metrics $\varepsilon_F, \varepsilon_Q$ used to monitor progress are the *normalized mean square root error* (nmse) from the corresponding projections \mathbf{z} :

$$\varepsilon_{\mathbf{a}}(\mathbf{z}) = \frac{\| [P_{\mathbf{a}} - I]\mathbf{z} \|}{\|\mathbf{a}\|},$$

$$\varepsilon_{\mathbf{Q}}(\mathbf{z}) = \frac{\| [P_{\mathbf{Q}} - I]\mathbf{z} \|}{\|\mathbf{a}\|}$$

where I is the identity operator. This has to be compared to ε_0 , the error w.r.t the known solution $\hat{\psi}$:

$$\begin{aligned} \varepsilon_0(\mathbf{z}) &= \frac{1}{\|\mathbf{a}\|} \min_{\varphi} \left\| e^{i\varphi} \mathbf{z} - \mathbf{Q}\hat{\psi} \right\|, \\ \varepsilon'_0(\mathbf{z}) &= \frac{1}{\|\mathbf{Q}^*\mathbf{Q}\hat{\psi}\|} \min_{\varphi} \left\| e^{i\varphi} \mathbf{Q}^*\mathbf{z} - \mathbf{Q}^*\mathbf{Q}\hat{\psi} \right\|, \\ \varepsilon'_0(\psi) &= \frac{1}{\|\mathbf{Q}^*\mathbf{Q}\hat{\psi}\|} \min_{\varphi} \left\| \mathbf{Q}^*\mathbf{Q}(e^{i\varphi}\psi - \hat{\psi}) \right\|, \end{aligned}$$

where φ is an arbitrary global phase factor.

The initial values for the input data and translations can either be loaded from file or set by a python interface. The starting “zero-th” initial image is either loaded from file, set to a random image, or taken as a constant

$\mathbf{Q}_{(i)}\psi$	$\mathbf{z}_{(i)}(\mathbf{r}) = w(\mathbf{r})\psi(\mathbf{r} + \mathbf{x}_{(i)})$
$\psi_{(i)} = \mathbf{Q}_{(i)}^* \mathbf{z}_{(i)}$	$\psi(\mathbf{r} + \mathbf{x}_{(i)}) = w^*(\mathbf{r})\mathbf{z}_{(i)}(\mathbf{r})$
$\mathbf{Q}^* \mathbf{z}$	$\sum_{\mathbf{x}_k + \mathbf{r}_k = \mathbf{l}} \text{conj}(w(\mathbf{r}_k)) \mathbf{z}_{(i)}(\mathbf{r}_k)$
$e_l^T \mathbf{Q}^* \mathbf{Q} e_k$	$\sum_{\mathbf{x}_k + \mathbf{r}_k = \mathbf{l}} w(\mathbf{r}_k) ^2 \delta_{l,k}$
$P_{\mathbf{a}_{(i)}} \mathbf{z}_{(i)}$	$\sum_{\mathbf{q}} e^{-i\mathbf{q} \cdot \mathbf{r}} \frac{\sum_{\mathbf{r}} e^{i\mathbf{q} \cdot \mathbf{r}} \mathbf{z}_{(i)}(\mathbf{r})}{ \sum_{\mathbf{r}} e^{i\mathbf{q} \cdot \mathbf{r}} \mathbf{z}_{(i)}(\mathbf{r}) } \mathbf{a}_{(i)}$

TABLE I. Linear algebra notation. Here e_l is the unit $n \times 1$ vector with the l -th entry 1 and δ is the Kronecker delta. The division is understood as an element-wise operation. The operator $P_{\mathbf{a}_{(i)}}$ is defined in Eq.(6)

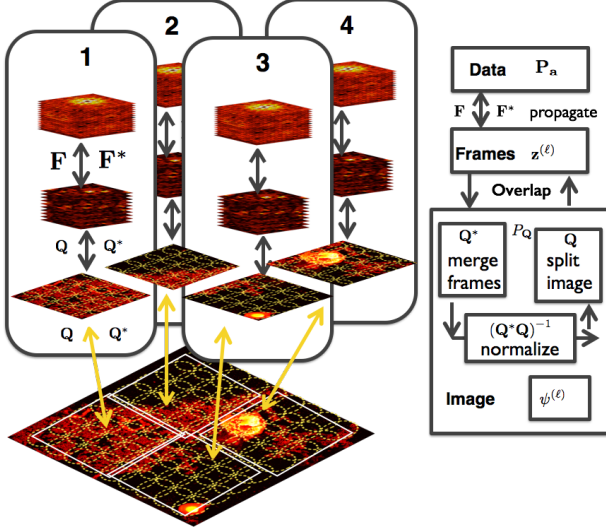


FIG. 2. To achieve the highest possible throughput and scalability one has to parallelize across multiple GPUs. As most ptychographic scans use a constant density of scan point across the object, we expect to be able to achieve a very even division, resulting in good load balancing. SHARP essentially enforces an overlap constraint between the images produced by each of the GPUs, and also enforces that the illumination recovered on each GPU agree with each other. This is done by default is done at every iteration.

image using flags. To achieve high performance, one has to compute the product of several linear operators (\mathbf{Q} , \mathbf{F} , \mathbf{S} , \mathbf{Q}^* , \mathbf{F}^* , \mathbf{S}^*) on a set of frames \mathbf{z} , an image ψ and an illumination w several times. We use a distributed GPU architecture across multiple nodes for this task (Fig. 2). Fourier transform operators \mathbf{F} , and \mathbf{F}^* are computationally intensive but also easily parallelized. To extract multiple frames out of an image ($\mathbf{Q}\psi$) we initially decompose the translations and corresponding frames onto multiple nodes in a regular grid. To add up all the frames onto an image ($\mathbf{Q}^* \mathbf{z}$) or updating the illumination \mathbf{S}^* require a reduction within each GPU and communication across nodes using MPI. The frequency of communication can be adjusted.

III. APPLICATIONS AND PERFORMANCE

SHARP enables high-throughput streaming analysis using computationally efficient phase retrieval algorithms. In this section we describe a typical dataset and sample that can be collected in less than 1 minute at the ALS, and the computational backend to provide fast feedback to the microscopist. To characterize our performance, we use both simulations and experimental data. We use simulations to compare convergence of the reconstruction algorithm to the “true solution” and characterize the effect of different light sources, contrast, scale, noise, detectors or samples for which no data exist yet. Experimental data from ALS used to characterize battery materials, green cement, at different wavelengths and orientation has been successfully reconstructed^{71–75} using the software described in this article. In a streaming example we describe a front-end that operates very close to the actual experiment, while the backends runs remotely on a GPU/CPU cluster.

A. Simulations and performance

As a demonstration, we start from a sample was composed of colloidal gold 50 nm and 10 nm nanoparticles deposited on a transparent silicon nitride membrane. An experimental image was obtained by scanning electron microscopy which provides high resolution and contrast but can only view the surface of the sample. We simulate a complex transmission function by scaling the image amplitude from 0 to 50 nm thickness, and using the complex index of refraction of Gold at 750 eV energy from [henke.lbl.gov]. The illumination is generated by zone-plate diameter of 220 microns, 60 nm outer zone width, discretized with (128×128) pixels in the far field.

Convergence to single precision numerical accuracy for a typical scan of 10,000 frames, is achieved in less than 2 seconds (see Fig. 3).

B. Experimental example

Figure 4 shows ptychographic reconstructions of a dataset generated from a sample consisting of gold balls with diameters of 50 and 10 nm. The data were generated using x-rays with energy of 750 eV at beamline 5.3.2.1 of the Advanced Light Source with exposure times of 1 second and a square scan grid with 40 nm spacing. The details of data acquisition and processing prior to reconstruction are reported elsewhere⁷⁶. The reconstructions consisted of 300 iterations of the RAAR algorithm with a probe retrieval and background retrieval step every other iteration while global synchronizations were performed every iteration. We note that both the phase and amplitude reconstruct with high signal-to-noise ratio, an advantage enabled by high stability soft x-ray scanning

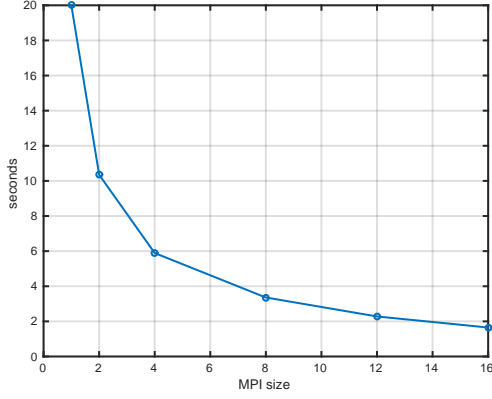


FIG. 3. Timing to process 10,000 frames of dimension 128×128 extracted from an image of 1000×1000 as a function of the number of nodes. Reconstruction ($\varepsilon'_0 < 5e - 4$) is achieved in under 2 seconds using a cluster of 4 nodes with 4 GTX Titan GPU per node. Similar performance can be expected in a single node using newer GPUs. Timing contributions for corresponding computational kernels are $(\mathbf{Q}^* \mathbf{Q})^{-1} \mathbf{Q}^*$ 30 %, \mathbf{F}, \mathbf{F}^* 20 %, \mathbf{Q} 20 %, \mathbf{S} 5 %, elementwise operations 20 %, and residual calculation 5 %.

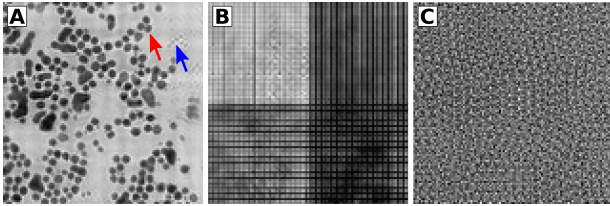


FIG. 4. Reconstruction of a test sample consisting of gold balls with diameters of 50 and 10 nm. A) Phase image generated by SHARP using the probe retrieval algorithm, background retrieval algorithm, and the Fourier mask applied to the probe. The red arrow points to a collection of 50 nm balls while the blue arrow points to a collection of 10 nm balls. The pixel size is 10 nm. B) Same as (A) except with the Fourier mask on the probe turned off. C) Same as (A) except with the background retrieval algorithm turned off.

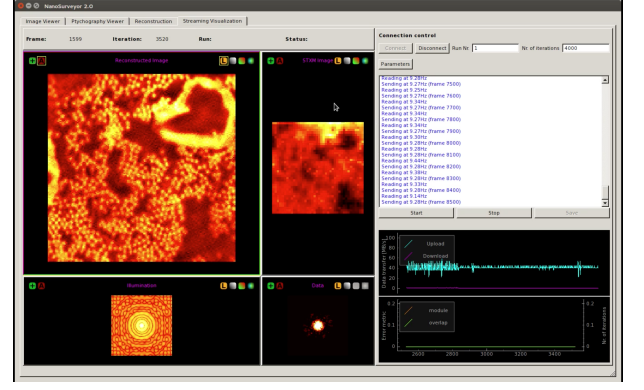


FIG. 5. Graphical User Interface (GUI) for the streaming pipeline. As soon as deata frames are recorded by the CCD camera, the different components in the pipeline are sending updates to the GUI.

transmission microscopes, and that features at the single pixel level are apparent.

C. Interface and Streaming

Common processing pipelines used for ptychographic experiments usually have a series of I/O operations and many different components involved. With the streaming pipeline, all components are connected, which allows users to monitor and quickly act upon changes along the experimental and computational pipeline.

A new run is triggered based on the user-defined settings for data aquisition and processing using a Graphical User Interface (Fig. 5). The GUI allows the user to stop and restart the run (including the streaming pipeline) at any given time, and also provides different saving schemes.

On the back-end side, the streaming infrastructure is divided into a handler of 4 different workers addressing dark frames, diffraction frames, reduced and down-sampled images and the ptychographic reconstruction (SHARP) as shown in Figure 6. The handler connects the back-end to the front-end and controls the communication and data flow among the different back-end workers which all have different roles. The dark worker is accumulating dark frames providing statistical maps (mean and variance) of the noise structure on the detector. Every frame worker is transforming raw into clean diffraction frames. This involves a subtraction of the average dark, filtering, photon counting and downsampling. Depending on the computing capacities of the back-end, it is possible to run N frame workers simultaneously. The image worker is reducing a collection of clean diffraction frames, producing low-resolution STXM images and an initial estimate of the illumination function. As soon as clean diffraction frames are ready and a guess for the illumination has been provided, the SHARP worker starts

- ²⁸ R. Hegerl and W. Hoppe, *Berichte Der Bunsen-Gesellschaft Fur Physikalische Chemie* **74**, 1148 (1970).
- ²⁹ P. D. Nellist, B. C. McCallum, and J. M. Rodenburg, *Nature* **374**, 630 (1995).
- ³⁰ H. N. Chapman, *Ultramicroscopy* **66**, 153 (1996).
- ³¹ J. M. Rodenburg, *Ptychography and Related Diffractive Imaging* (Elsevier, 2008) Chap. Ptychography and Related Diffractive Imaging Methods, pp. 87–184.
- ³² R. Hegerl and W. Hoppe, *Berichte Der Bunsen-Gesellschaft Fur Physikalische Chemie* **74**, 1148 (1970).
- ³³ J. M. Rodenburg and R. H. T. Bates, *Phil. Trans. R. Soc. Lond. A* **339**, 521 (1992).
- ³⁴ J. C. Spence, *High-resolution electron microscopy*, Vol. 60 (Clarendon Press, 2003).
- ³⁵ H. N. Chapman, *Ultramicroscopy* **66**, 153 (1996).
- ³⁶ P. Thibault, M. Dierolf, A. Menzel, O. Bunk, C. David, and F. Pfeiffer, *Science* **321**, 379 (2008).
- ³⁷ P. Thibault, M. Dierolf, O. Bunk, A. Menzel, and F. Pfeiffer, *Ultramicroscopy* **109**, 338 (2009).
- ³⁸ J. M. Rodenburg, A. C. Hurst, A. G. Cullis, B. R. Dobson, F. Pfeiffer, O. Bunk, C. David, K. Jefimovs, and I. Johnson, *Phys. Rev. Lett.* **98**, 034801 (2007).
- ³⁹ J. M. Rodenburg and H. M. L. Faulkner, *Appl. Phys. Lett.* **85**, 4795 (2004).
- ⁴⁰ C. Kewish, P. Thibault, M. Dierolf, O. Bunk, A. Menzel, J. Vila-Comamala, K. Jefimovs, and F. Pfeiffer, *Ultramicroscopy* **110**, 325 (2010).
- ⁴¹ S. Hönig, R. Hoppe, J. Patommel, A. Schropp, S. Stephan, S. Schöder, M. Burghammer, and C. G. Schroer, *Opt. Express* **19**, 16324 (2011).
- ⁴² M. Guizar-Sicairos, S. Narayanan, A. Stein, M. Metzler, A. R. Sandy, J. R. Fienup, and K. Evans-Lutterodt, *Applied Physics Letters* **98**, 111108 (2011).
- ⁴³ A. Maiden, M. Humphry, M. Sarahan, B. Kraus, and J. Rodenburg, *Ultramicroscopy* **120**, 64 (2012).
- ⁴⁴ M. Beckers, T. Senkbeil, T. Gorniak, K. Giewekemeyer, T. Salditt, and A. Rosenhahn, *Ultramicroscopy* **126**, 44 (2013).
- ⁴⁵ P. Thibault and M. Guizar-Sicairos, *New Journal of Physics* **14**, 063004 (2012).
- ⁴⁶ P. Godard, M. Allain, V. Chamard, and J. Rodenburg, *Opt. Express* **20**, 25914 (2012).
- ⁴⁷ N. C. Jesse and G. P. Andrew, *Applied Physics Letters* **99**, 154103 (2011).
- ⁴⁸ G. Zheng, R. Horstmeyer, and C. Yang, *Nature Photonics* **7**, 739 (2013).
- ⁴⁹ D. J. Batey, D. Claus, and J. M. Rodenburg, *Ultramicroscopy* **138**, 13 (2014).
- ⁵⁰ S. Dong, R. Shiradkar, P. Nanda, and G. Zheng, *Biomedical Optics Express* **5**, 1757 (2014).
- ⁵¹ J. Morrison, L. Rätty, P. Marriott, and P. O’Toole, *Scientific reports* **3** (2013).
- ⁵² L. Tian, X. Li, K. Ramchandran, and L. Waller, *Biomedical Optics Express* **5**, 2376 (2014).
- ⁵³ D. Vine, G. Williams, B. Abbey, M. Pfeifer, J. Clark, M. De Jonge, I. McNulty, A. Peele, and K. Nugent, *Physical Review A* **80**, 063823 (2009).
- ⁵⁴ M. Stockmar, P. Cloetens, I. Zanette, B. Enders, M. Dierolf, F. Pfeiffer, and P. Thibault, *Scientific reports* **3** (2013).
- ⁵⁵ S. Marchesini and H.-T. Wu, *Rank-1 accelerated illumination recovery in scanning diffractive imaging*, Method Rep. LBNL-6734E, arXiv:1105.5628 (Lawrence Berkeley National Laboratory, 2014) arXiv:1408.1922 [math.OC].
- ⁵⁶ A. Schropp, R. Hoppe, V. Meier, J. Patommel, F. Seiboth, H. J. Lee, B. Nagler, E. C. Galtier, B. Arnold, U. Zastrau, et al., *Scientific reports* **3** (2013).
- ⁵⁷ J. R. Fienup, *The Institute of Optics*, 1 (2003).
- ⁵⁸ Z. Wen, C. Yang, X. Liu, and S. Marchesini, *Inverse Problems* **28**, 115010 (2012).
- ⁵⁹ C. Yang, J. Qian, A. Schirotzek, F. Maia, and S. Marchesini, *Iterative Algorithms for Ptychographic Phase Retrieval*, Tech. Rep. 4598E, arXiv:1105.5628 (Lawrence Berkeley National Laboratory, 2011).
- ⁶⁰ S. Marchesini, A. Schirotzek, C. Yang, H.-t. Wu, and F. Maia, *Inverse Problems* **29**, 115009 (2013).
- ⁶¹ J. R. Fienup, J. C. Marron, T. J. Schulz, and J. H. Seldin, *Appl. Opt.* **32**, 1747 (1993).
- ⁶² B. Abbey, K. A. Nugent, G. J. Williams, J. N. Clark, A. G. Peele, M. A. Pfeifer, M. de Jonge, and I. McNulty, *Nature Physics* **4**, 394 (2008).
- ⁶³ L. W. Whitehead, G. J. Williams, H. M. Quiney, D. J. Vine, R. A. Dilanian, S. Flewett, K. A. Nugent, A. G. Peele, E. Balaur, and I. McNulty, *Phys. Rev. Lett.* **103**, 243902 (2009).
- ⁶⁴ M. Guizar-Sicairos and J. R. Fienup, *Opt. Express* **16**, 7264 (2008).
- ⁶⁵ S. T. Thurman and J. R. Fienup, *JOSA A* **26**, 1008 (2009).
- ⁶⁶ M. Guizar-Sicairos and J. R. Fienup, *Opt. Express* **17**, 2670 (2009).
- ⁶⁷ F. R. Maia, *Nature methods* **9**, 854 (2012).
- ⁶⁸ F. J. Anscombe, *Biometrika*, 246 (1948).
- ⁶⁹ M. Mäkitalo and A. Foi, *Image Processing, IEEE Transactions on* **22**, 91 (2013).
- ⁷⁰ R. Luke, *Inverse Problems* **21**, 37 (2005).
- ⁷¹ D. A. Shapiro, Y.-S. Yu, T. Tyliczszak, J. Cabana, R. Celestre, W. Chao, K. Kaznatcheev, A. Kilcoyne, F. Maia, S. Marchesini, et al., *Nature Photonics* (2014).
- ⁷² Y.-S. Yu, C. Kim, D. A. Shapiro, M. Farmand, D. Qian, T. Tyliczszak, A. D. Kilcoyne, R. Celestre, S. Marchesini, J. Joseph, et al., *Nano letters* **15**, 4282 (2015).
- ⁷³ S. Bae, R. Taylor, D. Shapiro, P. Denes, J. Joseph, R. Celestre, S. Marchesini, H. Padmore, T. Tyliczszak, T. Warwick, et al., *Journal of the American Ceramic Society* **98**, 4090 (2015).
- ⁷⁴ Y. Li, S. Meyer, J. Lim, S. C. Lee, W. E. Gent, S. Marchesini, H. Krishnan, T. Tyliczszak, D. Shapiro, A. L. D. Kilcoyne, et al., *Advanced Materials* **27**, 6591 (2015).
- ⁷⁵ Y. Li, S. Meyer, J. Lim, S. C. Lee, W. E. Gent, S. Marchesini, H. Krishnan, T. Tyliczszak, D. Shapiro, A. L. D. Kilcoyne, et al., *Advanced Materials* **27**, 6590 (2015).
- ⁷⁶ D. A. Shapiro, Y.-S. Yu, T. Tyliczszak, J. Cabana, R. Celestre, W. Chao, K. Kaznatcheev, K. L. David, F. Maia, S. Marchesini, Y. S. Meng, T. Warwick, L. L. Yang, and H. A. Padmore, *Nat Photon* **8**, 765 (2014).